arXiv:1601.05568v2 [stat.CO] 24 Feb 2016

EFFICIENT PARAMETER INFERENCE IN GENERAL HIDDEN MARKOV MODELS USING THE FILTER DERIVATIVES.

Jimmy Olsson and Johan Westerborn

KTH Royal Institute of Technology
Department of Mathematics
SE-100 44 Stockholm, Sweden

ABSTRACT

Estimating online the parameters of general state-space hidden Markov models is a topic of importance in many scientific and engineering disciplines. In this paper we present an online parameter estimation algorithm obtained by casting our recently proposed particle-based, rapid incremental smoother (PaRIS) into the framework of recursive maximum likelihood estimation for general hidden Markov models. Previous such particle implementations suffer from either quadratic complexity in the number of particles or from the well-known degeneracy of the genealogical particle paths. By using the computational efficient and numerically stable PaRIS algorithm for estimating the needed prediction filter derivatives we obtain a fast algorithm with a computational complexity that grows only linearly with the number of particles. The efficiency and stability of the proposed algorithm are illustrated in a simulation study.

***Index Terms*—** Hidden Markov models, maximum likelihood estimation, online parameter estimation, particle filters, recursive estimation, sequential Monte Carlo methods, state-space models

1. INTRODUCTION

This paper deals with the problem of *online parameter estimation* in general state-space *hidden Markov models* (HMMs) using *sequential Monte Carlo* (SMC) methods and a *recursive maximum likelihood* (RML) algorithm. HMMs with general state spaces, also referred to as *state-space models*, are currently applied within a large variety of scientific and engineering disciplines, see, e.g., [1, Chapter 1] and the references therein.

A *hidden Markov model* (HMM) is a bivariate model consisting of an observable process $\{Y_t\}_{t\in\mathbb{N}}$, known as the *observation process*, and an unobservable Markov chain $\{X_t\}_{t\in\mathbb{N}}$, known as the *state process*, taking values in some general state spaces Y and X, respectively. We let χ and q_θ denote the initial distribution and transition density (with respect to some reference measure denoted by $\mathrm{d}x$ for simplicity), respectively, of the hidden Markov chain, where $\theta \in \Theta$ denotes a parameter vector that fully determines the dynamics of model and Θ is the parameter space. Conditioned on the state process, the observations are assumed to be independent with conditional distribution of Y_t depending on X_t only. We denote by g_θ the density of the latter conditional distribution. Before using the HMM in practice, e.g. for prediction, the model needs to be calibrated through estimation of the parameters θ given observations $y_{0:t} = (y_0, \ldots, y_t)$ (this will be our standard notation for vectors). There are several different approaches to this problem, based on either frequentist or Bayesian inference, and [2] provides a good exposé of some current methods.

The *maximum likelihood* approach aims at finding the parameter vector that maximizes the *likelihood function* $\theta \mapsto L_\theta(y_{0:t})$ or, equivalently, the *log-likelihood function* $\theta \mapsto \ell_\theta(y_{0:t}) = \log L_\theta(y_{0:t})$. For the model specified above, the likelihood is given by

$$L_\theta(y_{0:t}) = \int g_\theta(x_0, y_0)\chi(x_0) \prod_{s=1}^{t} g_\theta(x_s, y_s) q_\theta(x_{s-1}, x_s)\, \mathrm{d}x_{0:t}.$$

The log-likelihood function and its gradient are generally intractable, as closed-form expressions are obtainable only in the linear Gaussian case or in the case where the state space X is a finite set. Still, maximization of the log-likelihood function can be performed using the following *Robbins-Monro scheme*: at iteration n, let $\theta_n = \theta_{n-1} + \gamma_n Z_n$, where Z_n is a noisy measurement of $\nabla\ell_{\theta_{n-1}}(y_{0:t})$, i.e. the gradient of the log-likelihood with respect to θ evaluated at $\theta = \theta_{n-1}$, and the sequence $\{\gamma_n\}_{n\in\mathbb{N}^*}$ of positive step sizes satisfies the regular stochastic approximation requirements $\sum_{n=1}^{\infty}\gamma_n = \infty$ and $\sum_{n=1}^{\infty}\gamma_n^2 < \infty$. Note that this approach requires the approximation Z_n to be recomputed at every iteration of the algorithm. However, if the number of observations is very large, computing Z_n is costly. Thus, since many iterations may generally be required for convergence, this yields an impractical algorithm.

This work is supported by the Swedish Research Council, Grant 2011-5577.

Hence, rather than incorporating the full observation record into each parameter update, it is preferable to update the parameters little by little as the data is processed. This is of course of particular importance in *online* applications where the observations become available "on-the-fly". One such scheme is the following RML approach, where we instead update iteratively the parameters according to

$$\theta_{t+1} = \theta_t + \gamma_{t+1}\zeta_{t+1}, \tag{1.1}$$

where ζ_{t+1} is a noisy observation of $\nabla \ell_{\theta_t}(y_{t+1} \mid y_{0:t})$, i.e. the gradient (with respect to θ), evaluated at $\theta = \theta_t$, of the log-density of Y_{t+1} given $Y_{0:t}$. This technique was, in the case of a finite state space X, studied in [3]. The same work also establishes, under certain conditions, the convergence of the output $\{\theta_t\}_{t\in\mathbb{N}}$ towards the parameter θ^* generating the data.

In this note we present an efficient particle implementation of the previous RML scheme based on the *particle-based, rapid incremental smoother* (PaRIS) introduced in [4] and analyzed further in [5]. Compared with previous implementations, which suffer typically from quadratic complexity in the number of particles [6, 7], our new algorithm is significantly faster as it has a computational complexity that grows only *linearly* in the particle population size. The rest of the paper is organized as follows. After having introduced some basic notation in Section 2, we delve deeper into the RML algorithm in Section 3. In Section 4 we present our new algorithm and finally we present, in Section 5, a simulation study, comparing our algorithm with an existing approach.

2. SOME NOTATION

In the following we let $\mathbb{N}$ denote the set natural numbers and set $\mathbb{N}^* = \mathbb{N} \setminus \{0\}$.

The conditional distribution of $X_{s:s'}$ given fixed observations $Y_{0:t} = y_{0:t}$ (with $(s, s') \in \mathbb{N}^2$ and $s \le s' \le t$) is given by

$$\phi_{s:s'|t;\theta}(x_{s:s'}) = L_\theta^{-1}(y_{0:t}) \iint g_\theta(x_0, y_0)\chi(x_0) \times \prod_{u=1}^{t} g_\theta(x_u, y_u) q_\theta(x_{u-1}, x_u) \,\mathrm{d}x_{0:s-1} \,\mathrm{d}x_{s'+1:t}.$$

We refer to $\phi_{t;\theta} = \phi_{t|t;\theta}$ and $\phi_{0:t|t;\theta}$ as the *filter* and *joint smoothing distributions* at time t, respectively. The distribution

$$\phi_{t+1|t;\theta}(x_{t+1}) = \int q_\theta(x_t, x_{t+1})\phi_{t+1|t;\theta}(x_t) \,\mathrm{d}x_t,$$

i.e., the conditional distribution of X_{t+1} given $Y_{0:t} = y_{0:t}$ is referred to as the *prediction filter* at time t. In general, these distributions are intractable and need to be approximated. We will often consider expectations of some integrable function f with respect to the posterior above and write $\phi_{s:s'|t;\theta}(f) = \int f(x_{s:s'})\phi_{s:s'|t;\theta}(x_{s:s'}) \,\mathrm{d}x_{s:s'}$.

For a vector of non-negative numbers $\{a_\ell\}_{i=1}^N$ we let $\mathsf{Pr}(\{a_\ell\}_{i=1}^N)$ denote the *categorical distribution* induced by $\{a_\ell\}_{\ell=1}^N$. In other words, $J \sim \mathsf{Pr}(\{a_\ell\}_{\ell=1}^N)$ means that the random variable J takes on the value i with probability $a_i / \sum_{\ell=1}^N a_\ell$.

3. RECURSIVE MAXIMUM LIKELIHOOD

As mentioned in the introduction, the RML approach updates the parameter for each new observation y_{t+1} using (1.1). The algorithm that we propose is based on the fact that $\nabla \ell_{\theta_t}(y_{t+1} \mid y_{0:t})$ can be decomposed as

$$\begin{aligned}
&\nabla \ell_{\theta_t}(y_{t+1} \mid y_{0:t}) \\
&= \bigg(\int \phi_{t+1|t;\theta_t}(x_{t+1}) \nabla g_{\theta_t}(x_{t+1}, y_{t+1}) \,\mathrm{d}x_{t+1} \\
&+ \int \nabla \phi_{t+1|t;\theta_t}(x_{t+1}) g_{\theta_t}(x_{t+1}, y_{t+1}) \,\mathrm{d}x_{t+1} \bigg) \\
&\times \bigg(\int \phi_{t+1|t;\theta_t}(x_{t+1}) g_{\theta_t}(x_{t+1}, y_{t+1}) \,\mathrm{d}x_{t+1} \bigg)^{-1},
\end{aligned} \tag{3.1}$$

where $\nabla \phi_{t+1|t;\theta_t}(x_{t+1})$ is known as the *prediction filter derivative* [6, 7] or *tangent filter*. In order to compute the second integral in (3.1) we follow the lines of [7, Section 2]; more specifically, write, for some $\phi_{t+1|t;\theta}$-integrable function f,

$$\begin{aligned}
&\int \nabla \phi_{t+1|t;\theta}(x_{t+1}) f(x_{t+1}) \,\mathrm{d}x_{t+1} \\
&\qquad = \mathbb{E}[f(X_{t+1}) h_{t+1}(X_{0:t+1}; \theta) \mid y_{0:t}] \\
&\qquad - \mathbb{E}[f(X_{t+1}) \mid y_{0:t}] \mathbb{E}[h_{t+1}(X_{0:t+1}; \theta) \mid y_{0:t}],
\end{aligned} \tag{3.2}$$

where $h_t(x_{0:t}; \theta) = \sum_{s=0}^{t-1} \tilde{h}_s(x_{s:s+1}; \theta)$, with

$$\tilde{h}_s(x_{s:s+1}; \theta) = \begin{cases} \nabla \log \{g_\theta(x_s, y_s) q_\theta(x_s, x_{s+1})\} & \text{for } s \in \mathbb{N}^*, \\ \nabla \log g_\theta(x_0, y_0) & \text{for } s = 0, \end{cases}$$

is of additive form. In addition, using the tower property of conditional expectations, we may elaborate further (3.2) according to

$$\begin{aligned}
&\int \nabla \phi_{t+1|t;\theta}(x_{t+1}) f(x_{t+1}) \,\mathrm{d}x_{t+1} \\
&\qquad = \int \phi_{t+1|t;\theta}(x_{t+1}) \bigg\{ \tau_{t+1}(x_{t+1}; \theta) \\
&- \int \tau_{t+1}(x_{t+1}; \theta) \phi_{t+1|t;\theta}(x_{t+1}) \,\mathrm{d}x_{t+1} \bigg\} f(x_{t+1}) \,\mathrm{d}x_{t+1},
\end{aligned} \tag{3.3}$$

with

$$\tau_{t+1}(x_{t+1}; \theta) = \mathbb{E}[h_{t+1}(X_{0:t+1}; \theta) \mid X_{t+1} = x_{t+1}, y_{0:t}]. \tag{3.4}$$

Combining the updating scheme (1.1) with (3.1) and (3.3), with $f(x_{t+1})$ playing the role of $g_{\theta_t}(x_{t+1}, y_{t+1})$ in (3.3), gives us an RML algorithm, in which the noisy measurements $\{\zeta_t\}_{t\in\mathbb{N}^*}$ are produced via (3.1) and (3.3) using a particle-based approach approximating online the prediction filter distributions as well as the statistics (3.4).

4. PARIS-BASED RML

A *particle filter* updates sequentially, using importance sampling and resampling techniques, a set $\{(\xi_t^i, \omega_t^i)\}_{i=1}^N$ of particles and associated weights targeting a sequence of distributions. In our case we will use the particle filter to target the prediction filter flow $\{\phi_{t+1|t;\theta}\}_{t\in\mathbb{N}}$ in the sense that for all $t \in \mathbb{N}$ and $\phi_{t+1|t;\theta}$-integrable functions f,

$$\phi_{t+1|t;\theta}^N(f) = \frac{1}{N}\sum_{i=1}^N f(\xi_{t+1}^i) \backsimeq \phi_{t+1|t;\theta}(f) \quad \text{as } N \to \infty.$$

While particle filters are generally well-suited for approximation of filter and prediction filter flows, estimation of the sequence $\{\tau_t\}_{t\in\mathbb{N}}$, where each τ_t is a conditional expectation over a *path space*, is a considerably more challenging task. As established in several papers (see, e.g., [5, 8, 9]), standard sequential Monte Carlo methods fall generally short when used for path-space approximation due to particle ancestral lineage degeneracy. Since $h_t(x_{0:t};\theta)$ is additive, we express however the sequence $\{\tau_t\}_{t\in\mathbb{N}}$ *recursively* as

$$\begin{aligned}&\tau_{t+1}(x_{t+1};\theta)\\ &= \mathbb{E}[\tau_t(X_t;\theta) + \tilde{h}_t(X_{t:t+1};\theta) \mid X_{t+1} = x_{t+1}, y_{0:t}].\end{aligned} \quad (4.1)$$

Our aim is to approximate the recursion (4.1) on a grid of particles. Thus, letting, for $i \in \{1,\dots,N\}$, $\tilde{\tau}_t^i$ be approximations of $\tau_t(\xi_t^i;\theta)$, we may approximate the recursion (4.1) by

$$\tilde{\tau}_{t+1}^i = \sum_{j=1}^N \frac{\omega_t^j q_\theta(\xi_t^j, \xi_{t+1}^i)}{\sum_{\ell=1}^N \omega_t^\ell q_\theta(\xi_t^\ell, \xi_{t+1}^i)} \left(\tilde{\tau}_t^j + \tilde{h}_t(\xi_t^j, \xi_{t+1}^i;\theta)\right), \quad (4.2)$$

where the ratio serves as a particle approximation of the *backward kernel*, i.e., the conditional distribution of X_t given $y_{0:t}$ and $X_{t+1} = \xi_{t+1}^i$; see [9, 10]. The recursion (4.2) is the key component of the *particle RML* proposed in [7, Alg. 2]. However, since each update requires a sum of N terms to be computed, this approach has a computational complexity that grows quadratically with the the number of particles N.

As an alternative, the PaRIS algorithm [5] provides an efficient and numerically stable estimator of smoothed expectations of additive form and hence a more convenient way of approximating the updating formula (4.1). More specifically, we replace each update (4.2) by the PaRIS-type update

$$\tau_{t+1}^i = \tilde{N}^{-1}\sum_{j=1}^{\tilde{N}} \left(\tau_t^{J_{t+1}^{(i,j)}} + \tilde{h}_t(\xi_t^{J_{t+1}^{(i,j)}}, \xi_{t+1}^i;\theta)\right),$$

where $\tilde{N} \in \mathbb{N}^*$ is typically small and $\{J_{t+1}^{(i,j)}\}_{j=1}^{\tilde{N}}$ are i.i.d. indices drawn from $\mathsf{Pr}(\{\omega_t^\ell q(\xi_t^\ell, \xi_t^i)\}_{\ell=1}^N)$. Using an accept-reject sampling approach proposed in [9], each index can be drawn efficiently with a cost that can be proven to be uniformly bounded in N and t. This yields an algorithm with complexity $\mathcal{O}(N\tilde{N})$ instead of $\mathcal{O}(N^2)$. A key discovery made in [5] is that the PaRIS algorithm converges, as N tends to infinity, for all *fixed* $\tilde{N} \in \mathbb{N}^*$ and stays numerically stable, as t tends to infinity, when $\tilde{N} \geq 2$; in particular, it is, for large N and when $\tilde{N} \geq 2$, possible to derive a time-uniform $\mathcal{O}(\{1 + 1/(\tilde{N}-1)\}/N)$ bound on the variance of the PaRIS estimator, suggesting that $\tilde{N}$ should be kept at a *moderate value* (as increasing $\tilde{N}$ will not lead to significant reduction of variance). In fact, using only $\tilde{N} = 2$ draws has turned out to work well in simulations (see [5, Section 4]). We refer to [5] for a detailed discussion. Compared to the $\mathcal{O}(N^2)$ algorithm based on the updates (4.2), the extra randomness introduced by the PaRIS-based algorithm adds some variance to the estimates; still, it will be clear in the simulations that follow that the linear computational complexity of our approach allows considerably more particles to be used, resulting in a significantly smaller variance for a given computational budget.

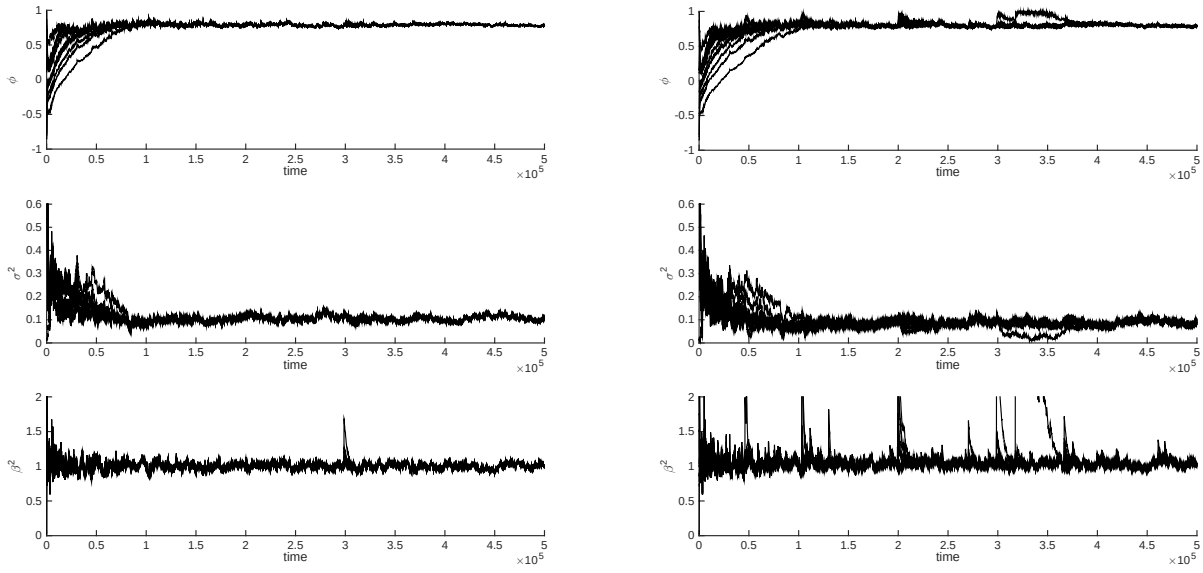
4.1. PaRIS-based RML algorithm

In the algorithm below, all operations involving the index i should be performed for all $i \in \{1,\dots,N\}$. The precision parameter $\tilde{N}$ is to be set by the user. In addition ζ_t^1, ζ_t^2 and ζ_t^3 correspond to the three terms of (3.1).

Initialize the algorithm by setting arbitrarily θ_0
Draw $\xi_0^i \sim \chi$
Set $\tau_0^i \leftarrow 0$
for $t \leftarrow 0, 1, 2, \dots$ **do**
  Set $\omega_t^i \leftarrow g_{\theta_t}(\xi_t^i, y_t)$
  Draw $I_i \sim \mathsf{Pr}(\{\omega_t^i\}_{i=1}^N)$
  Draw $\xi_{t+1}^i \sim q_{\theta_t}(\xi_t^{I_i}, \cdot)$
  for $j \leftarrow 1, \dots, \tilde{N}$ **do**
    Draw $J_{t+1}^{(i,j)} \sim \mathsf{Pr}(\{\omega_t^\ell q_{\theta_t}(\xi_t^\ell, \xi_{t+1}^i)\}_{\ell=1}^N)$
  end for
  Set

$$\tau_{t+1}^i \leftarrow \tilde{N}^{-1}\sum_{j=1}^{\tilde{N}} \left(\tau_t^{J_{t+1}^{(i,j)}} + \tilde{h}_t(\xi_t^{J_{t+1}^{(i,j)}}, \xi_{t+1}^i;\theta_t)\right)$$

  Set $\bar{\tau}_{t+1} \leftarrow N^{-1}\sum_{\ell=1}^N \tau_{t+1}^\ell$
  Set $\zeta_{t+1}^1 \leftarrow N^{-1}\sum_{\ell=1}^N \nabla g_{\theta_t}(\xi_{t+1}^\ell, y_{t+1})$
  Set $\zeta_{t+1}^2 \leftarrow N^{-1}\sum_{\ell=1}^N \left(\tau_{t+1}^\ell - \bar{\tau}_{t+1}\right) g_{\theta_t}(\xi_{t+1}^\ell, y_{t+1})$
  Set $\zeta_{t+1}^3 \leftarrow N^{-1}\sum_{\ell=1}^N g_{\theta_t}(\xi_{t+1}^\ell, y_{t+1})$
  Set $\theta_{t+1} \leftarrow \theta_t + \gamma_{t+1}\dfrac{\zeta_{t+1}^1 + \zeta_{t+1}^2}{\zeta_{t+1}^3}$
end for



(a) PaRIS-based RML

(b) Particle RML

Fig. 1. Particle learning trajectories produced by the PaRIS-based (left panel) and standard particle (right-panel) RML for, from top to bottom, ϕ, σ^2 and β^2. For each algorithm, 12 replicates were generated on the same data set with different, randomized initial parameters (being the same for both algorithms). For the particle RML, the plot of β^2 does not contain the full trajectories due to very high peaks.

5. SIMULATIONS

We tested our method on the stochastic volatility model

$$\begin{aligned} X_{t+1} &= \phi X_t + \sigma V_{t+1}, \\ Y_t &= \beta \exp(X_t/2) U_t, \end{aligned} \qquad t \in \mathbb{N},$$

where $\{V_t\}_{t\in\mathbb{N}^*}$ and $\{U_t\}_{t\in\mathbb{N}}$ are independent sequences of mutually independent standard Gaussian noise variables. Parameters to be estimated were $\theta = (\phi, \sigma^2, \beta^2)$, and we compared the performance of our PaRIS-based RML to that of the particle RML proposed in [7]. To get a fair comparison of the algorithms we set the number of particles used in each algorithm such that both algorithms ran in the same computational time. With our implementation, $N = 100$ for the particle RML corresponded to $N = 1400$ and $\tilde{N} = 2$ for the PaRIS-based RML. For both algorithms we set $\gamma_t = t^{-0.6}$. The algorithms were executed on data comprising $500\,000$ observations generated under the parameter $\theta^* = (0.8, 0.1, 1)$. Each algorithm ran 12 times with the same observation input but with randomized starting parameters (still, the same starting parameters were used for both algorithms). In Fig. 1 we present the resulting learning trajectories, and it can clearly be seen that the PaRIS-based RML exhibits significantly less variance in its estimates, especially for the β^2 variable. In the particle RML we notice some large jumps in the β^2 variable, originating from the fact that the corresponding estimate of ζ_t^3 gets very small. This is due to the low number of particles failing to cover the support of the emission density. In contrast, since we can utilize considerably more particles in the PaRIS-based RML, we see only a single, comparably small, jump in β^2. Judging by the estimated (on the basis of the 12 trajectories) variances $(.054, .164, .063) \times 10^{-3}$ and $(.069, .181, .095) \times 10^{-4}$ of the final parameter estimates for the particle RML and the PaRIS-based RML, respectively, the PaRIS-based RML is roughly ten times more precise than the particle RML.

6. DISCUSSION

We have proposed a novel algorithm for online parameter learning in general HMMs using an RML method based on the PaRIS algorithm [5]. The new method has a linear computational complexity in the number of particles, which allows considerably more particles to be used for a given computational budget compared to previous methods. The performance of the algorithm is illustrated by simulations indicating clearly improved convergence properties of the parameter estimates.

7. REFERENCES

[1] O. Cappé, E. Moulines, and T. Rydén, *Inference in Hidden Markov Models*, Springer, 2005.

[2] N. Kantas, A. Doucet, S. S. Singh, J. Maciejowski, and N. Chopin, "On particle methods for parameter estimation in state-space models," *Statist. Sci.*, vol. 30, no. 3, pp. 328–351, 08 2015.

[3] F. Le Gland and L. Mevel, "Recursive estimation in HMMs," in *Proc. IEEE Conf. Decis. Control*, 1997, pp. 3468–3473.

[4] J. Olsson and J. Westerborn, "Efficient particle-based online smoothing in general hidden Markov models," in *IEEE 2014 International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2014)*, 2014.

[5] J. Olsson and J. Westerborn, "Efficient particle-based online smoothing in general hidden Markov models: the PaRIS algorithm," *Bernoulli*, 2016, to appear.

[6] G. Poyiadjis, A. Doucet, and S. S. Singh, "Particle methods for optimal filter derivative: application to parameter estimation," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 18-23 March 2005, pp. v/925–v/928.

[7] P. Del Moral, A. Doucet, and S. S. Singh, "Uniform stability of a particle approximation of the optimal filter derivative," *SIAM Journal on Control and Optimization*, vol. 53, no. 3, pp. 1278–1304, 2015.

[8] J. Olsson, O. Cappé, R. Douc, and E. Moulines, "Sequential Monte Carlo smoothing with application to parameter estimation in non-linear state space models," *Bernoulli*, vol. 14, no. 1, pp. 155–179, 2008.

[9] R. Douc, A. Garivier, E. Moulines, and J. Olsson, "Sequential Monte Carlo smoothing for general state space hidden Markov models," *Ann. Appl. Probab.*, vol. 21, no. 6, pp. 2109–2145, 2011.

[10] P. Del Moral, A. Doucet, and S. Singh, "Forward smoothing using sequential Monte Carlo," Tech. Rep., Cambridge University, 2010.